

**SHARP:  
Automated Monitoring of Spacecraft Health and Status\***

David J. Atkinson, Mark L. James, R. Gaius Martin

Jet Propulsion Laboratory, California Institute of Technology  
4800 Oak Grove Drive, Pasadena, CA 91109

**ABSTRACT**

This paper briefly describes the spacecraft and ground systems monitoring process at the Jet Propulsion Laboratory and highlights some difficulties associated with the existing technology used in mission operations. A new automated system based on artificial intelligence technology is described which seeks to overcome many of these limitations. The system, called the Spacecraft Health Automated Reasoning Prototype (SHARP), is designed to automate health and status analysis for multi-mission spacecraft and ground data systems operations. The SHARP system has proved to be effective for detecting and analyzing potential spacecraft and ground systems problems by performing real-time analysis of spacecraft and ground data systems engineering telemetry. Telecommunications link analysis of the Voyager 2 spacecraft was the initial focus for evaluation of the system in a real-time operations setting during the Voyager spacecraft encounter with Neptune in August, 1989. The SHARP system will be delivered to the JPL Space Flight Operations Center for regular use by planetary flight projects, including the Galileo and Magellan spacecraft, and will also be applied to monitoring and control applications in the Deep Space Network's Network Operations Control Center.

**2. INTRODUCTION**

The Voyager 1 and Voyager 2 spacecraft were launched from Cape Canaveral, Florida, on August 20, 1977. The technology to monitor the health and status of these probes was designed and developed in the early 1970's. This now-antiquated technology, coupled with the heroic efforts of many JPL personnel over the last 13 years, has carried Voyager 2

---

\* This paper was previously presented at SPIE, Orlando, Florida, April, 1990.

through near-fatal catastrophic events to four of our solar systems outer planets. Despite the spacecraft's failed radio receiver, sunlight damage to the photopolarimeter scientific instrument, and partially paralyzed scan platform (which houses Voyager's imaging system), JPL engineers have kept Voyager operational, enabling the capture and transmission of vast amounts of invaluable information and images of the Jovian, Saturnian, Uranian, and Neptunian systems.

During critical periods of the mission, up to 40 real-time operators are required to monitor the spacecraft's 10 subsystems on a 24-hour, 7-day-per-week schedule. This does not include the numerous subsystem and scientific instrument specialists who must constantly be available on call to handle emergencies. Unlike the 1980's, when JPL mission operations could focus on the two Voyager spacecraft, in the coming decade there will be an increasing number of planetary exploration spacecraft flying at the same time. In addition to the Voyagers, the Galileo and Magellan spacecraft have been launched in the past year and are now on their way to Jupiter and Venus, respectively. The Ulysses, CRAF (Comet Rendezvous and Asteroid Flyby), Mars Observer, and other spacecraft will follow in the next few years. To accommodate the increasing load on mission operations, JPL has established a Space Flight Operations Center (SFOC) to replace the individual mission control teams and spacecraft teams for each mission. A single, multi-mission flight team will operate all of the spacecraft. As more spacecraft are launched and begin to carry out their missions, the Space Flight Operations Center will require significant advances in automation technology in order to support the increasing workload on operations personnel and to ensure the safety of the spacecraft.

The Spacecraft Health Automated Reasoning Prototype (SHARP) was developed as part of an on-going effort to apply artificial intelligence (AI) techniques to mission operations automation. The primary task for an operational SHARP system will be multi-mission monitoring and diagnosis of spacecraft and ground systems in the Space Flight Operations Center. As tools such as SHARP are developed, they are demonstrated and evaluated in tough, operational settings to prove their performance. The Voyager 2 spacecraft was targeted for the initial demonstration of the SHARP system. The spacecraft's August 1989 encounter with the planet Neptune afforded an excellent opportunity to evaluate SHARP in a rigorous environment. The monitoring and troubleshooting of the telecommunications subsystem on-board Voyager 2 and the process of real-time telecommunications link analysis were selected as the initial operations functions to be automated. Telecommunications with the

Voyager 2 spacecraft suffers from frequent anomalies and requires coordination of monitoring and diagnosis efforts of both the spacecraft and ground telecommunications systems. Due to cumbersome and time-consuming manual processes and obsolete technology which will be discussed in later paragraphs, severe limitations exist on the current methods of analyzing Voyager telecommunications data. Even with the substantial improvement in computing support which is part of the new Space Flight Operations Center, the telecommunications area is both an operations area sorely in need of automation as well as one of the most challenging to automate.

### **3. TELECOMMUNICATIONS OPERATIONS**

This section gives a brief overview of the telecommunications mission operations process, specifically focusing on the monitoring of spacecraft telecommunications subsystem health and telecommunications link status operations. Two of the major challenges for automation are described: The automation of manual data processing and data interpretation, and the automated real-time anomaly detection and analysis.

As noted earlier, each spacecraft is monitored on a continuous basis. To enable the receipt and collection of spacecraft engineering data, JPL operates three complexes of antennas located around the world. These complexes comprise NASA's Deep Space Network (DSN). With the exception of occultations and a short gap between two of the stations (Canberra and Madrid), a spacecraft is always in view from one of these Deep Space Stations (DSS), as the complexes are called. A scheduled observing period for a station is called a pass.

Three of the most important functions which are part of analysis of the telecommunications link between the spacecraft, Deep Space Network, and ground system computers at JPL are, 1) the numerical estimation of telecommunications subsystem and link performance, 2) the monitoring of real-time telecommunications activity and detection of failures or degraded performance, and 3) the diagnosis, isolation, and recovery from these problems. To accomplish each of these functions, a wide variety of information must be accessed and processed manually by an operator.

Predictions of telecommunications performance are embodied in a type of data known as "Predicts". Predicts are precise, numerical estimations of expected engineering data values for particular spacecraft and Deep Space Station parameters that impact the performance of the telecommunications link, such as signal-to-noise ratio and antenna elevation. Predicts are generated for each spacecraft pass over each

ground station and can be divided into four categories: raw predictions, pass predictions, instantaneous predictions, and residual calculations. While the details of Predict generation and analysis are beyond the scope of this paper, it can be noted that much of the Predict calculation process is performed manually, and is tedious, time-consuming, incomplete, and error-prone. Telecommunications operators may spend up to two hours each day computing Predicts by hand using hardcopied listings of spacecraft activity, raw predictions, and pocket calculators. The SHARP system completely automates the process of Predict generation and analysis, saving up to two hours of operator time each day.

In addition to Predicts, telecommunications operators use the "Integrated Sequence of Events" (ISOE) to aid in monitoring telecommunications activity. The ISOE is a hardcopy listing of scheduled spacecraft and Deep Space Network activity. Operators use the ISOE in Predict calculations, alarm determination, and anomaly diagnosis. The operators must visually scan the ISOE to highlight relevant telecommunications information. This process is prone to error during periods of high spacecraft activity and when operators unknowingly do not reference the latest activity modifications to the ISOE. The SHARP system maintains a current, on-line database of ISOE information and automatically provides relevant telecommunication activity information from the ISOE to the system's other real-time monitoring processes and the operator as needed.

The monitoring of telecommunications and detection of anomalies is further complicated by the selection of alarm limits for spacecraft and Deep Space Station engineering parameters. Unlike the Predict values which are precise numerical predictions arising from a quantitative simulation of spacecraft performance, the engineering alarm limits are critical thresholds which define the acceptable range of engineering values on any telemetry channel. Excursions beyond the alarm limit range indicate imminent failure situations. In current Voyager spacecraft operations, alarm limits are determined manually according to the information in the ISOE, design information about spacecraft subsystem performance, and "rules of thumb" arising from the spacecraft team's experience with actual subsystem performance over the life of a mission. The current manual procedure to change alarm limits is so impeditive that for many engineering data channels typically a wide threshold is selected that incorporates the entire range of parameter conditions, thereby creating a risk of undetected anomalies. (See Doyle<sup>1</sup> for a discussion of problems in the determination of alarm limits).

In telecommunications as in other areas, the ultimate diagnosis, isolation, and recovery from failures, anomalous conditions, or degraded system performance often requires the intervention of experts who have years of specialized experience operating spacecraft subsystems (e.g., power, thermal, telecommunications). One of the most serious limitations on the current method of mission operations are the critical flight skills built up by these experts over the many years of flying spacecraft. These specialists must be on-call at any time, and are frequently consulted on a daily basis. The timeliness of an expert response to a problem can be critical in saving a spacecraft. Furthermore, when the experts retire, their critical skills are lost to mission operations. The Voyager 2 spacecraft has already been flying for almost 13 years, and is expected to operate until 2018. Many future spacecraft are expected to have similar longevity. The accumulated expertise of mission operations personnel is a critical resource which should be preserved, and not recreated every time a senior engineer leaves the flight project.

#### **4. DESCRIPTION OF THE SHARP SYSTEM**

The SHARP system applies artificial intelligence as well as conventional computer science techniques to automate and eliminate much of the tedious data processing and analysis associated with the monitoring of spacecraft and ground system health and status. Many of the manual, labor-intensive and error prone activities are eliminated in part or whole by SHARP. Some of these were described in the previous section. The major automated functions provided by the SHARP system include:

- Real-time anomaly detection and diagnosis;
- Visualization of channelized data and system status;
- Acquisition and centralization of engineering data in a single workstation;
- Real-time analysis of spacecraft performance predictions;
- Integration with specialized numerical analysis software, e.g., Fast Fourier Transforms for determining spacecraft antenna pointing accuracy.

Figure 1 illustrates a top-level view of the SHARP system. Shown are the individual modules that comprise the system, as well as relevant components that are external to the Voyager application of SHARP. SHARP is implemented in Common LISP on a Symbolics 3650 color LISP Machine. The system is currently being ported to a Sun workstation, also running Common LISP. SHARP relies extensively on an expert system building language called STAR\*TOOL, developed at JPL<sup>2</sup>. The remainder of this

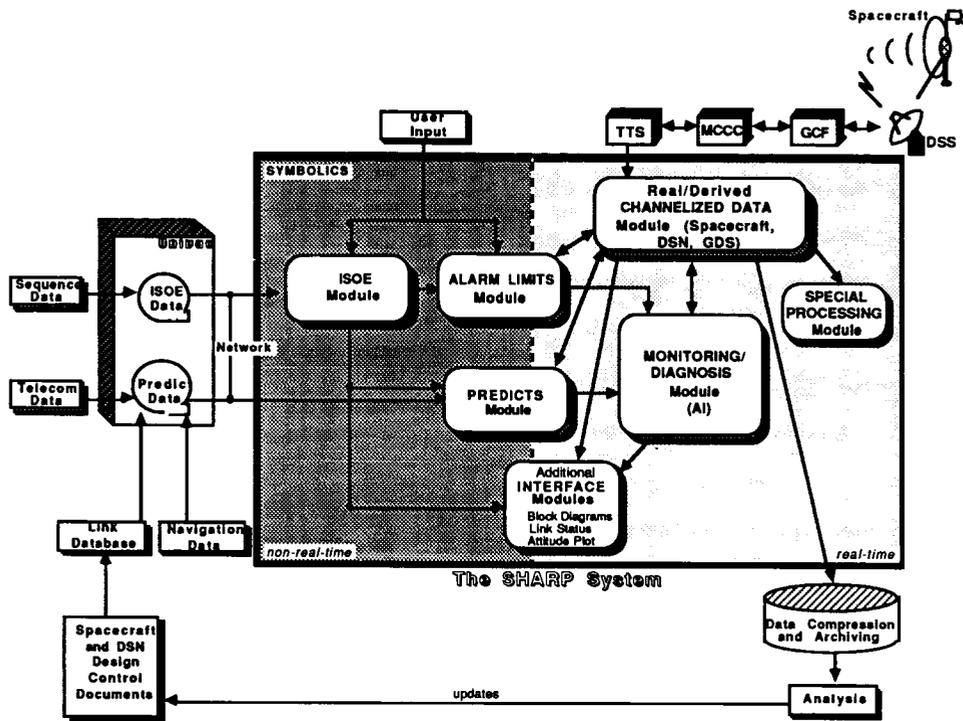


Figure 1. SHARP Telecom System Overview

paper will focus on the first of these automated functions: real-time anomaly detection and diagnosis. The remaining SHARP functions are described elsewhere<sup>3</sup>.

In SHARP, the automation of fault detection and diagnosis is accomplished through the use of artificial intelligence programming techniques. Artificial intelligence techniques are distributed throughout all components of the SHARP system. Artificial intelligence programming methodologies have enabled more effective automation and thorough analysis for SHARP functions. Unlike the current manual methods used in space flight operations, fault detection and diagnosis in SHARP is extremely fast, taking approximately 1/200th of a second from receipt of anomalous data to determination of a diagnosis. This speed is directly attributable to the AI techniques incorporated by the design of the system. Some of the techniques used in the SHARP system include: Procedural reasoning, blackboards, reasoning using context trees, heuristic adaptive parsing, and spontaneous computation daemons. Figure 2 illustrates the

ORIGINAL PAGE IS  
OF POOR QUALITY

design of the "AI Module" in SHARP, which is responsible for fault detection and diagnosis.

#### **4.1 Alarm determination**

The first step in verifying nominal spacecraft performance is to determine whether received engineering data values are within acceptable limits. Data which is outside limits is considered in alarm, and must be explained. Data values can be classified as nominal, in "soft alarm" (possibly indicating a warning condition), and in "hard alarm" (possibly indicating an imminent failure condition). SHARP makes this determination automatically, by selecting the appropriate alarm limits for each channel of data and comparing new data against those limits in real-time.

The SHARP module responsible for this function is the Alarm Executive, as shown in Figure 2. The Alarm Executive module has a predetermined model of spacecraft states and transitions between those states. The Voyager application of SHARP has 39 such states. Alarm limits on each engineering data channel are determined in advance by the domain expert for each one of these spacecraft states. The limits are represented in table format, and organized hierarchically into a discrimination network 7 layers deep (the network is ultimately compiled into a very efficient internal representation).

When a new engineering datum is received, the Alarm Executive first scans the Integrated Sequence of Events (ISOE) for the major activities and specialized activities which determine the spacecraft's current state, and further confirms the state by checking real-time engineering data related to spacecraft configuration. These are the keys used to search the spacecraft state discrimination network. In the case of the Voyager application of SHARP, the automatic gain control lock is checked to see if it is synchronized. The correct table of alarm limits is retrieved and the datum is matched against the appropriate alarm limits within the table after any additional conditions are checked, such as operator overrides. In general, more than a simple comparison of the datum against minimum and maximum threshold values is possible in determining an alarm condition. For example, an arbitrary function can be invoked to determine whether an alarm condition exists. These functions can break down the engineering datum into its individual bit status for example, or look at derivative information for trend detection.

In some cases, an anomalous spacecraft condition is directly indicated, e.g., based on error codes in the engineering data. In most cases, however,

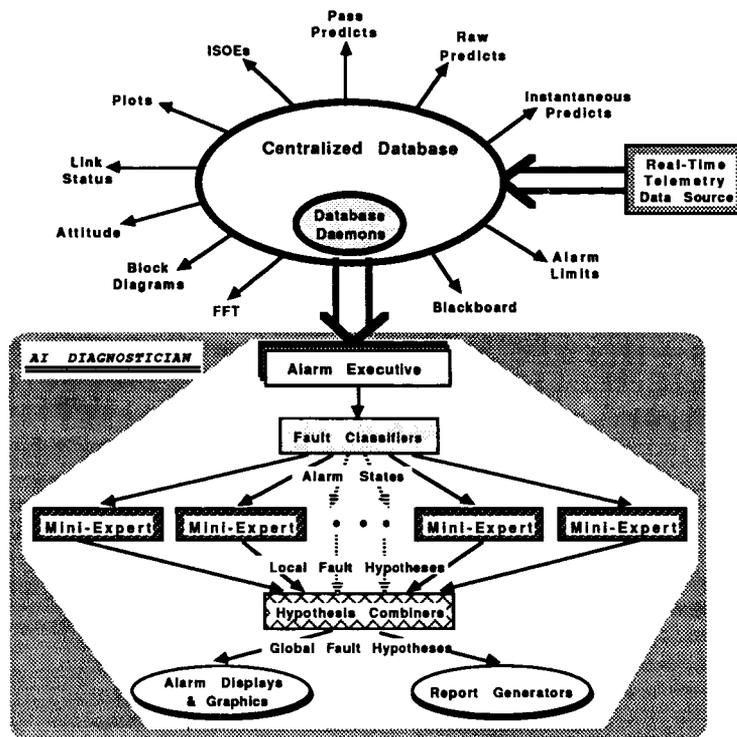


Figure 2. SHARP Artificial Intelligence Module

further analysis is required in order to determine the nature of the problem. The Alarm Executive makes this decision, and in addition monitors, logs, and reports to the telecommunications operator a number of attributes of the alarm situation, including the severity of alarm changes (i.e., from soft to hard alarm), the previous alarm status of the channel, and whether the operator has acknowledged the previous alarm messages. A variety of user interface and display changes are triggered by the Alarm Executive. If further analysis is required, the Alarm Executive informs the Fault Classifier module in SHARP. Analysis of alarm conditions by the Alarm Executive and Fault Classifier modules can proceed in parallel for any number of detected alarms.

#### 4.2 Fault Classification

The Fault Classification module is a rule-based system which makes an initial interpretation of alarm conditions, spacecraft state, and the sources of engineering data indicating the anomaly. The result of the

interpretation is a rough classification of the type of problem or its possible location in the telecommunications system, e.g., is it a spacecraft receiver problem, a possible configuration mismatch between the ground and spacecraft telecommunications subsystems, and so on. Frequently, there is no unambiguous interpretation available and subsequent diagnosis must proceed in parallel with several conflicting hypotheses. The products of the fault classification are asserted into a database which results in a pattern-directed invocation of specialized diagnostic routines, called "Mini-experts", described below. This architecture of hierarchical invocation of specialized diagnostic knowledge is related to the paradigm of cooperating specialists in classificatory diagnosis embodied in the CSRL system<sup>4</sup> and in the StarPlan system<sup>5</sup>.

#### **4.3 Mini-expert diagnostic routines**

The Voyager telecommunications application of SHARP includes approximately 40 mini-experts. These specialized diagnostic routines are each responsible for the local diagnosis of a specific fault or class of faults, such as particular channels in alarm, conical scan errors, configuration mismatches, or loss of telemetry. Mini-experts can be either cooperating or non-cooperating. A non-cooperating mini-expert focuses only on its designated fault area, and generally its conclusions can and should be reported independently to the operator. A cooperating mini-expert has the additional capability of searching beyond its local area to identify related faults that are likely to occur. In the process of this search, the cooperating mini-expert triggers other mini-experts who are specialists in those related areas. Information is exchanged between the mini-experts using a blackboard message system.

Mini-experts encode a procedural network of diagnostic decisions and analyses. They are related to rules in the Procedural Reasoning System (PRS) of Georgeff and Lansky<sup>6</sup>, although the representation mini-expert procedures differs. Mini-expert rule definitions include high-level descriptions of preconditions, activation and execution contexts, spacecraft state descriptions, relevant real-time data sources, hypotheses, and sequences of analyses and decisions which are part of the diagnostic process. Mini-expert knowledge definitions are not interpreted by SHARP. Instead, SHARP contains a compiler which generates Common LISP code from mini-expert descriptions and automatically installs the definitions into the SHARP run-time environment. The compiler performs the necessary bookkeeping and also checks for consistency with the other, installed mini-experts. Currently, a trained knowledge engineer must

develop mini-expert definitions by hand, and this constitutes a bottleneck for application of the system. To aid in knowledge acquisition, we are developing a graphical interface, called a "visual rule-building system" which can be used by domain experts to create mini-experts which would then be directly compiled by SHARP as before.

As mentioned above, the Fault Classification module may not determine a unique mini-expert to invoke. In this case, multiple mini-experts are invoked which pursue diagnoses in pseudo-parallel. Pseudo-parallelism is implemented in SHARP using facilities provided by STAR\*TOOL, which includes parallelism as a fundamental control structure. The various mini-experts and their rules operate in isolation of one another by executing in independent contexts<sup>7</sup> provided in the STAR\*TOOL memory model. Contexts can be organized into a tree-like structure to represent contradictory information resulting from changes in facts or from the introduction of new or contradictory hypotheses.

#### **4.4 Hypothesis Combination**

The Hypothesis Combiner module has the role of combining multiple fault hypotheses generated when several mini-experts are invoked in parallel by the Fault Classification module. The module communicates with mini-experts through SHARP's blackboard. Related fault hypotheses are combined into a single, more encompassing explanation for the operator (e.g., when there is a single action to take in response). Redundant hypotheses are eliminated in the process as well. When there are conflicting explanations for a detected problem, SHARP presents all of the explanations to the operator along with the separate recovery recommendations. In some cases, the operator is privy to information and knowledge which SHARP does not have, and can effectively disambiguate the situation. In any event, the final problem determination step and any corrective actions is left to the operator in cases of ambiguity.

### **5. VOYAGER ENCOUNTER WITH NEPTUNE EVALUATION**

Approximately one month before the Voyager encounter with the planet Neptune, a Symbolics workstation with SHARP loaded on it was moved from the Artificial Intelligence Laboratory at JPL into the real-time telecommunications operations area for the Voyager spacecraft. There were severe restrictions on how SHARP could interact with other Voyager systems. To simplify the installation, SHARP obtained spacecraft engineering data from the Voyager Test and Telemetry System over a printer port. Unabridged Integrated Sequence of Events and raw Predict

data were loaded into SHARP using tapes, rather than through network connections as in the Artificial Intelligence Laboratory.

During the demonstration period, SHARP helped find the cause of a Voyager science data error anomaly which appeared in the telemetry from the spacecraft as an excess error count. The SHARP system's graphical displays were used by telecommunications personnel to identify the problem and to characterize its magnitude. The problem was isolated using SHARP and other, manual trouble-shooting techniques to the Voyager ground data system and was corrected by the replacement of a wide-band interface unit in the Voyager Data Acquisition and Capture System (DACS). SHARP helped verify that the replacement of the unit actually fixed the problem. In a matter of hours, SHARP was able to assist operators in solving an anomalous condition which could have easily escalated to a more serious problem during the encounter itself, and could have taken human operators days or weeks to isolate without SHARP.

Also during the demonstration period, the knowledge engineer of SHARP and the domain expert would review alarms that SHARP had given. Generally, these alarms were correct. In one alarm situation, SHARP was giving warnings about the loss of the telecommunications signal. This ultimately turned out to be a false alarm as the spacecraft was undertaking a particular maneuver that the SHARP knowledge base did not contain, thereby leading the diagnostic system into an erroneous conclusion about antenna pointing. In other cases, SHARP was able to detect conditions where the Deep Space Station antenna tracking the spacecraft was drifting off point. SHARP detected these problems in a matter of seconds, and reported the condition to the telecommunications operators. Unfortunately, due to their previous lack of ability to detect and diagnose antenna pointing problems, the real-time telecommunications operators at JPL did not have procedures for alerting the Deep Space Station operators (possibly on the other side of the world) to antenna drift situations detected by SHARP. When the antenna drift reached a sufficient magnitude and urgency for the station operators to notice and correct, SHARP was able to detect the resolution of the problem and cancel the alarm situation. SHARP detected and correctly diagnosed other non-critical problems with the receiver automatic gain control and the S-band travelling wave tube temperature on board the spacecraft.

On the whole, the encounter with Neptune went extremely smoothly for the Voyager spacecraft. SHARP did not get a chance to make any really dramatic diagnoses, and the diagnostic system described in this paper did

not get a strenuous operational test. This underscores the difficulty in testing the diagnostic ability of real-time monitoring and control expert systems in operation settings: you may not get any problems! Using simulated data (based on historical problems with the spacecraft and based on synthetic situations) we were able to test SHARP much more thoroughly in the laboratory. SHARP is able to analyze 39 classes of telecommunications problems, and make about 60 unique diagnoses which require some problem-solving by the mini-experts to determine. Another 20 telecommunications problems are detectable by SHARP, but can be reported directly to the operator. Our domain expert estimates that SHARP covers approximately 80% of the known types of faults experienced in spacecraft telecommunications for Voyager. The remaining 20% include diagnoses which could be made if SHARP had the appropriate real-time data and additional knowledge engineering. As with most complex systems, there is always the possibility of novel faults. SHARP does not have the ability to successfully diagnose and explain a novel type of fault (nor was it intended to), but we are confident in the system's ability to detect departures from expected, nominal behavior.

## **6. EXPECTED BENEFITS FROM SHARP APPLICATIONS**

There are four principle areas where the JPL telecommunications users of SHARP expect to see benefits from application of the system and its descendents, which we are now developing. These areas are safety, workforce savings, reliability, and productivity.

Through its accurate detection, analysis, and tracking of the antenna drift and pointing conditions during the encounter, SHARP showed that it can detect and analyze important problems in a matter of seconds which currently take human operators minutes or hours. This provides an extra margin for ensuring the safety of the spacecraft, and thereby supports the success of the mission as a whole. The SHARP Voyager telecommunications domain expert, a man with over 20 years of experience who has cognizance not only for Voyager telecommunications operations but for other spacecraft as well, as stated publicly that the Soviets would not have lost the first Phobos spacecraft if they had SHARP applied to their telecommunications. One of the stated causes of the loss of the Phobos spacecraft has been that the spacecraft antenna drifted until the telecommunications link was lost due to a faulty attitude control command.

A second major benefit from application of SHARP will be in the area of workforce savings. Through its automation of many manual functions, SHARP promises to reduce the real-time link analysis operations staff by

a factor of five, and there is reason to believe that similar savings may be possible in other operations areas. This is precisely the type of benefit from automation which is necessary to support the single multi-mission flight team in the new JPL Space Flight Operations Center.

The system-wide status monitoring afforded by SHARP, and not discussed in detail in this paper, helps operators assure correct telecommunications system configuration. This is expected to reduce the number of commanding errors to the spacecraft and ground systems, and thereby reduce the loss or corruption of data due to configuration problems.

Finally, the SHARP system is expected to enhance the productivity of operations personnel by freeing them from the tedium of watching raw data and interpreting it for themselves. SHARP shifts the burden of routine monitoring operations, and most of the boring, manual computations which are involved, away from the operator to itself. This will enable operations personnel to perform required analyses more efficiently, and to exert a higher level of "supervisory monitoring" over multiple spacecraft subsystems on multiple spacecraft.

## **7. CONCLUSIONS**

Spacecraft and ground data systems operations present a rigorous environment in the area of monitoring and anomaly detection and diagnosis. With a number of planetary missions scheduled for the near future, the effort to staff and support these operations will present significant challenges.

The SHARP system was developed to address the challenges of automation in a multi-mission operations environment by augmenting conventional automation technologies with artificial intelligence. Its successful development and demonstration have led to a number of important conclusions. First and foremost, artificial intelligence technology is ready for application to spaceflight operations. The techniques can be used alongside conventional computer science techniques, and diagnostic knowledge-based systems can be embedded in the resulting application system. Acceptable real-time performance can be achieved. SHARP was never pushed to the limit of its speed or memory resources; in fact, most of its time was spent idle, waiting for new engineering data to process. This gives us confidence for broadening the approach in SHARP to multiple spacecraft subsystems.

The evaluation by Voyager personnel also taught us that the types of automation provided by SHARP are highly desired by operations personnel,

and are not viewed as job-threatening (although they may be in some cases). Operators were able to readily use the system with minimal training, and were enthusiastic about using the wide variety of graphical displays and options.

SHARP is now being extended and developed to a higher level of readiness so that flight projects such as Voyager, Magellan, Galileo, and others can use it directly. The system will be completed in 1990 and delivered to the Space Flight Operations Center for further evaluation and application to Magellan telecommunications. Separately, SHARP is also being applied to the Deep Space Network, Network Operations Control Center at JPL, with an operational system planned for 1991. Applications for remote monitoring and control of spaceborne instruments and experiments are also under consideration.

## **8. ACKNOWLEDGEMENTS**

The research described in this publication was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. The authors wish to acknowledge Denise Lawson and Harry Porta for their contributions to SHARP development, and to Boyd Madsen for providing expert knowledge for Voyager telecommunications. Portions of this work have been reported previously while in progress at the 1989 Goddard Conference on Space Applications of Artificial Intelligence, Greenbelt, Maryland, May 16-17.

## **9. REFERENCES**

1. Doyle, R., Sellers, S., and Atkinson, D., "Enhancing Aerospace System Autonomy Through Predictive Monitoring", *Proceedings of AIAA Meeting on Aerospace Sciences*, American Institute for Aeronautics and Astronautics, Reno, Nevada, 1989.
2. James, Mark, and Atkinson, David, "STAR\*TOOL -- An Environment and Language for Expert System Implementation", *Jet Propulsion Laboratory Report NTR C-17536*, August, 1988.
3. Lawson, D., and James, M., "SHARP: A Multi-Mission Artificial Intelligence System for Spacecraft Telemetry Monitoring and Diagnosis", *1989 Goddard Conference on Space Applications of Artificial Intelligence*, pp. 185-200, NASA, Greenbelt, Maryland, 1989. Also *Jet Propulsion Laboratory Publication 89-23*, May, 1989.

4. Chandrasekaran, B., "Distributed Knowledge-based Systems for Diagnosis and Information Retrieval", *Department of Computer and Information Sciences Report 763180/714659*, November, 1983.
5. Siemens, R.W., Golden, M., and Ferguson, J., "StarPlan II: Evolution of an Expert System", *Proceedings of the Fifth National Conference on Artificial Intelligence*, American Association for Artificial Intelligence, Philadelphia, PA, 1986.
6. Georgeff, M., and Lansky, A., "A System for Reasoning in Dynamic Domains: Fault Diagnosis on the Space Shuttle", *Technical Note 375*, Artificial Intelligence Center, SRI International, January, 1986.
7. McDermott, D. V., "Very Large PLANNER-type data bases", *Memo AIM-339*, AI Laboratory, Massachusetts Institute of Technology, 1975.